

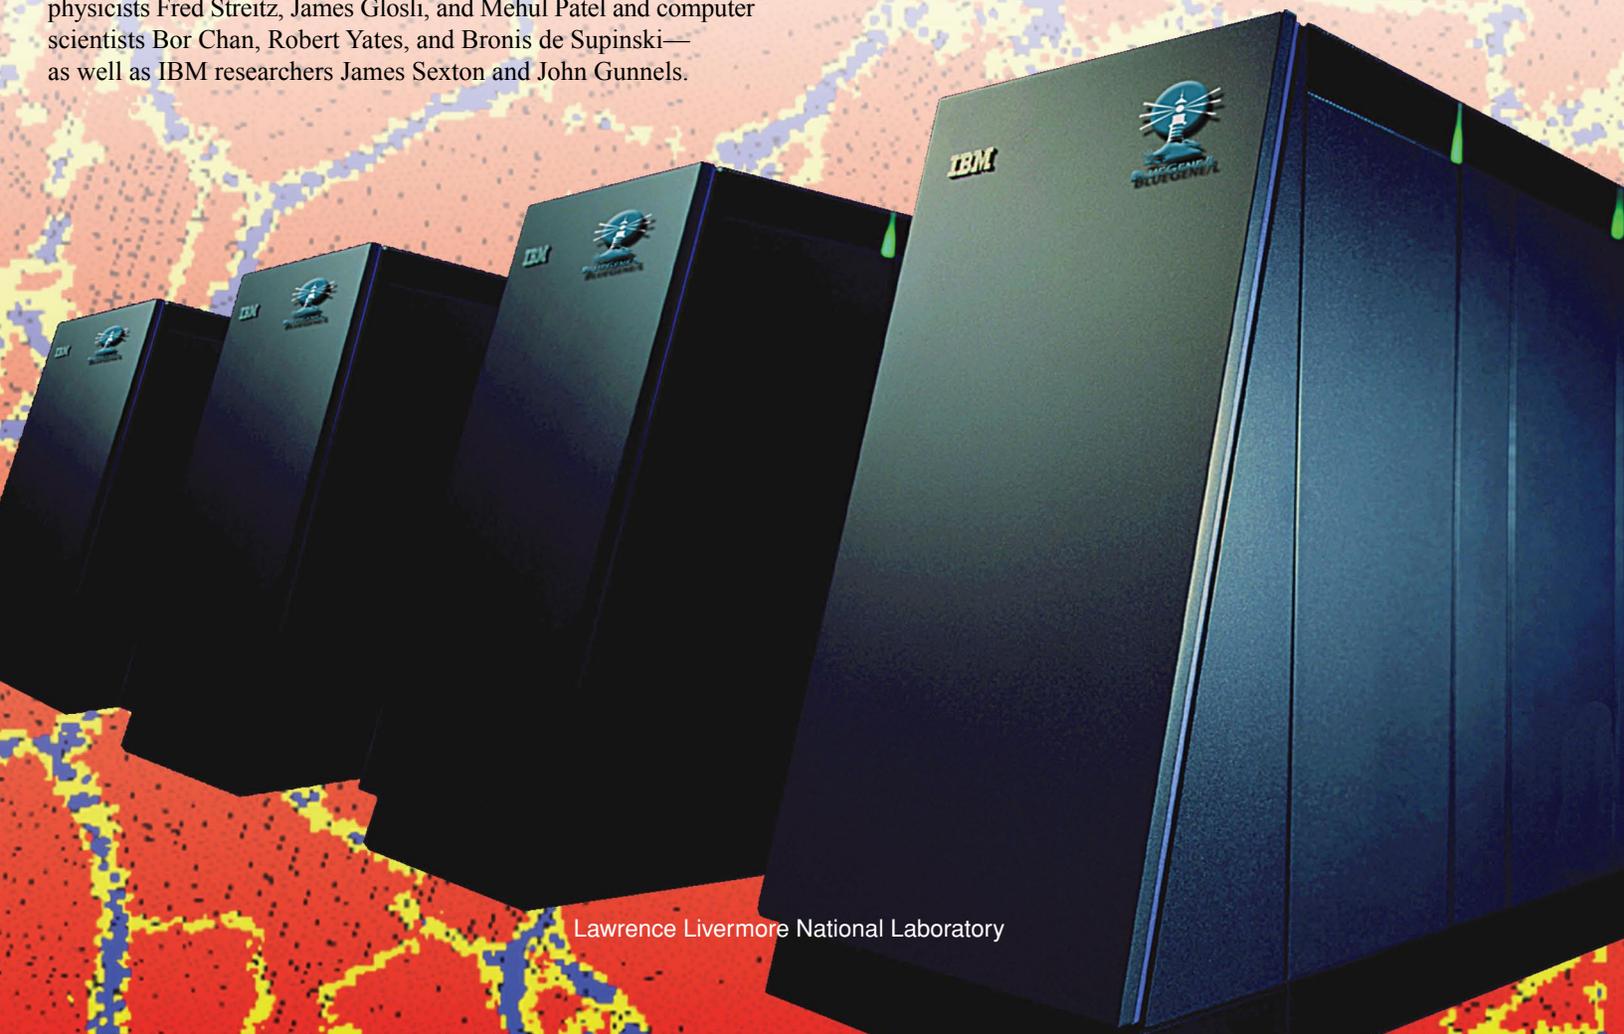
Keeping an Eye on the Prize

SETTING performance goals is part of the business plan for almost every company. The same is true in the world of supercomputers. Ten years ago, the Department of Energy (DOE) launched the Accelerated Strategic Computing Initiative (ASCI) to help ensure the safety and reliability of the nation's nuclear weapons stockpile without nuclear testing. ASCI, which is now called the Advanced Simulation and Computing (ASC) Program and is managed by DOE's National Nuclear Security Administration (NNSA), set an initial 10-year goal to obtain computers that could process up to 100 trillion floating-point operations per second (teraflops). Many computer experts thought the goal was overly ambitious, but the program's results have proved them wrong.

Last November, a Livermore-IBM team received the 2005 Gordon Bell Prize for achieving more than 100 teraflops while modeling the pressure-induced solidification of molten metal. The prestigious prize, which is named for a founding father of supercomputing, is awarded each year at the Supercomputing Conference to innovators who advance high-performance computing. Recipients for the 2005 prize included six Livermore scientists—physicists Fred Streitz, James Glosli, and Mehul Patel and computer scientists Bor Chan, Robert Yates, and Bronis de Supinski—as well as IBM researchers James Sexton and John Gunnels.

This team produced the first atomic-scale model of metal solidification from the liquid phase with results that were independent of system size. The record-setting calculation used Livermore's domain decomposition molecular-dynamics (ddcMD) code running on BlueGene/L, a supercomputer developed by IBM in partnership with the ASC Program. BlueGene/L reached 280.6 teraflops on the Linpack benchmark, the industry standard used to measure computing speed. As a result, it ranks first on the list of Top500 Supercomputer Sites released in November 2005.

To evaluate the performance of nuclear weapons systems, scientists must understand how materials behave under extreme conditions. Because experiments at high pressures and temperatures are often difficult or impossible to conduct, scientists rely on computer models that have been validated with obtainable data. Of particular interest to weapons scientists is the solidification of metals. "To predict the performance of aging nuclear weapons, we need detailed information on a material's phase transitions," says Streitz, who leads the Livermore-IBM team. For example, scientists



want to know what happens to a metal as it changes from molten liquid to a solid and how that transition affects the material's characteristics, such as its strength.

A New Code for Complex Systems

One metal the team simulated was the transition metal tantalum. In transition metals, the valence electrons, which interact with other elements to form compounds, are present in more than one shell. Thus, as tantalum solidifies, complex bonding structures form, and the transition from a melt phase to a solid can happen very slowly. These physical processes are challenging to model. A simulation of tantalum solidification may require billions of atoms, and the code must run many millions of time steps even though the process being simulated may last no more than a few nanoseconds.

Researchers have been modeling systems with billions of atoms for about 10 years. However, these models rely mainly on pair-potential techniques to describe the force each atom exerts on every other atom. Pair-potential techniques are effective for simple systems, such as those involving noble gases. Because noble gases have closed shells of electrons, the forces exerted on the atom are radially symmetric, resulting in spherically symmetric bonds.

Pair-potential techniques do not model complex systems with the accuracy needed for stockpile stewardship research. Most of the transition metals—including tantalum—contain a partially filled d band of electrons, which results in a more complicated bonding structure. For example, forces exerted on atoms are angularly dependent, and bonds may form between three or four atoms in a surrounding area. Accurately modeling these forces requires a sophisticated interaction potential.

In 1990, Livermore physicist John Moriarty developed the model-generalized pseudo-potential theory (MGPT), which can be used to

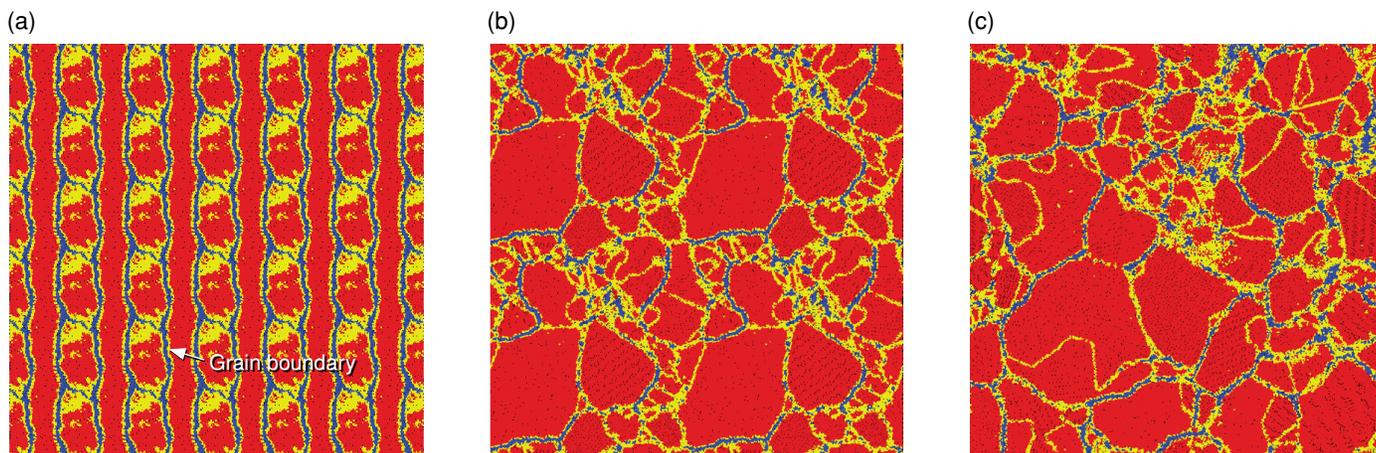
derive more accurate quantum-based interaction potentials. MGPT potentials are based on many-body expansions of a quantum-mechanically derived energy surface and include terms for two-, three-, and four-atom bonds. These potentials are validated by comparing information obtained by first-principles calculations and experiments.

Streitz and Patel first used MGPT potentials in 2000 in a single-processor code they had developed to model metal solidification. In 2002, Glosli joined the team, restructuring the MGPT potential routine and increasing single-processor performance by a factor of 20. He quickly became the principal architect of what is now the ddcMD code, leading the design and implementation of a novel domain decomposition algorithm that enabled parallel processing.

In the first full machine run on BlueGene/L, the team clocked the ddcMD code at 75 teraflops, significantly close to the ASC goal of 100 teraflops on a production science code. By focusing on the small matrix–matrix multiplication routines at the heart of the MGPT potentials, Glosli, Chan, and Gunnels boosted performance on short benchmark simulations to more than 107 teraflops. During a 7-hour production science run using all 131,072 processors, the team measured ddcMD performance at 101.7 teraflops—the highest sustained performance of a scientific application code.

Integrated Design a Major Advantage

The major difference between BlueGene/L and other computers is its scalability, which is provided by a large number of low-power processors and multiple integrated interconnection networks. BlueGene/L has 65,536 nodes, compared with 512 nodes for Livermore's ASC White machine, 1,536 for Purple, and 2,048 for the Q machine at Los Alamos National Laboratory. To accommodate this large number of nodes, IBM designed BlueGene/L with a simple



The Livermore–IBM team used three system sizes to simulate the transition of tantalum from a melt to a solid. (a) With a 64,000-atom system, the simulation produced an artificial grain-boundary pattern. (b) A simulation with more than 2 million atoms yielded a more realistic but restricted grain-size distribution. (c) The result from using 16 million atoms showed grain formation and growth independent of system size.

architecture that includes only 10 chips per node: nine memory chips and one compute application-specific integrated circuit (ASIC) chip. In comparison, a desktop computer can have 50 to 60 individual chips. The ASIC is a complete system-on-a-chip. It includes two IBM PowerPC 440 processors and five interconnects, and it provides 8 megabytes of embedded dynamic random access memory.

BlueGene/L's highly integrated design scales up in an orderly fashion with relatively modest power and cooling requirements. In 2004, when BlueGene/L assumed the number 1 spot on the Top500 list, it did so with only one-quarter of the final system, clocking 70.72 teraflops. (See *S&TR*, April 2005, pp. 23–25.) BlueGene/L beat the previous record holder, Japan's Earth Simulator, by a factor of two.

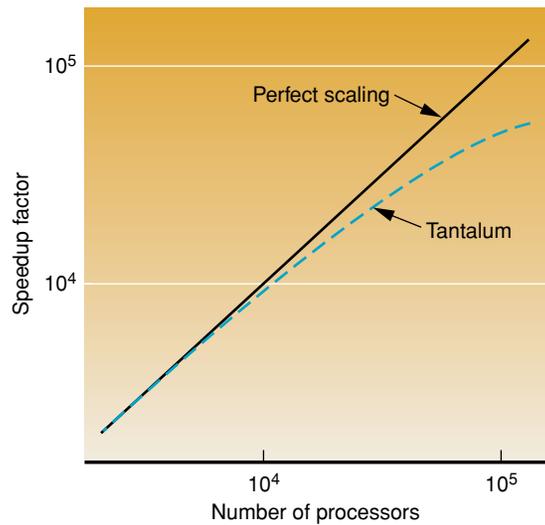
With the ddcMD code running on BlueGene/L, weapons scientists can simulate billions of atoms on the necessary time scales to obtain reliable results. "Prior to BlueGene/L, the added computational expense of the potentials we needed for the tantalum studies would have limited us to about 10,000 atoms for a 1-nanosecond simulation," says Streitz. "Not only would that calculation have taken a month to run, but the system being simulated would still be about 20 times smaller in size than what we need to access the required physics."

For the molten tantalum studies, the team modeled systems ranging from 64,000 to 128 million atoms compressed to 250 gigapascals of pressure at 5,000 kelvins. By varying the size of a simulation, the researchers gained confidence that their results were not affected by the size of the system being modeled. The 64,000-atom simulation showed two large grains, with a grain boundary spanning the simulation cell—an unrealistic result that indicated the system size was too small. In the simulation with more than 2 million atoms, the distribution of grain sizes was much more realistic. When the model size reached 16 million atoms, grain formation and growth were completely independent of system size. These simulations are the first step toward modeling nucleation—when the transition to the solid phase begins—and growth in a manner that allows scientists to directly link processes at atomistic scales to those at micrometer scales and above.

Doesn't Skip a Beat with Added Load

With many high-performance codes, increasing the number of processors will eventually slow computing performance substantially because the processors need more time to communicate with each other. In contrast, the ddcMD code can achieve excellent scaling performance on BlueGene/L because of the algorithms that Glosli incorporated into the code. It maintains almost perfect scalability even as the number of processors is increased from less than 1,000 to more than 100,000.

This scalability also allows researchers to adjust the system size to extract specific information. "Although a 64,000-atom system is too small for modeling molten tantalum through to its solid phase, the nucleation event may be the same in the small system as it is in



The domain decomposition molecular-dynamics code demonstrates excellent scalability on BlueGene/L as the number of processors working on a 16-million-atom system is increased.

a 16-million-atom system," says Streitz. "We can glean valuable data even from smaller system sizes."

The team is sifting through the information produced by the solidification simulations on BlueGene/L. "We have a mammoth amount of data that we still need to go through," says Glosli. "We may find some surprises in the results."

Long-Term Goals Pay Off

The team's simulations will help scientists develop larger-scale models of material behavior. They also are providing more information about the nucleation and growth processes that occur during solidification and how factors such as temperature and strain rate affect these processes. "Ultimately," says Glosli, "we want to build models that reduce processor time even further."

Surpassing 100-plus teraflops using a scientific application marks an important milestone for supercomputing. The simulations of metal solidification are providing valuable insight for NNSA's stockpile stewardship efforts to ensure the safety and reliability of the nation's nuclear deterrent. The acknowledgment of this record-setting achievement by the Gordon Bell Prize demonstrates that BlueGene/L can deliver as promised.

—Gabriele Rennie

Key Words: Accelerated Strategic Computing Initiative (ASCI), Advanced Simulation and Computing (ASC) Program, BlueGene/L, domain decomposition molecular-dynamics (ddcMD) code, Gordon Bell Prize, model-generalized pseudo-potential theory (MGPT), solidification.

For further information contact Fred Streitz (925) 423-3236 (streitz1@llnl.gov).